

Synchronising cell transmission for packet switching

The invention relates to a method for synchronising start of cell times in input/output means with cell transmission periods in at least one cross-connection means for packet switching, where cells are transferred between said input/output means by said cross-connection means in cell transfer periods, where configurations of said cross-connection means are changed between cell transfer periods in cross-connection configuration periods, where cells from said input/output means are sent at start of cell times, and where said sent cells are received in said cross-connection means. The invention further relates to a packet switch comprising input/output means, with a port controller, with a cell input port, and a cell output port, cross-connection means comprising cell input ports and cell output ports connected to said cell output port and cell input port of said port controllers, respectively. Furthermore, the invention relates to the use of such a method and such a packet switch.

Switching-nodes of packet switched networks comprise packet-switches. These packet-switches transfer data packets between input and output ports, based on address information comprised in each incoming packet. Incoming packets are buffered in line cards and organised in input port queues. These input port queues are organised as virtual output queues (VOQ), which are implemented inside a port controller, located on a line card. These port controllers are connected to switch cards, comprising crossbar matrices and an arbiter, respectively, by cables. The arbiter calculates input/output configurations of the matrices to allow an even transmission of incoming packets of all connected port controllers. To switch the incoming packets to the respective output ports, the crossbar matrices connect input ports with the respective output ports during cell transfer periods. To change the connections between input and output ports, the configuration of the matrices have to be changed by the arbiter during switch-over times.

Incoming packets are segmented within the line cards into fixed size packet fragments, also called cells. Outgoing packets are reassembled from cells which have been switched from a crossbar matrix to the respective line card.

To allow a smooth transmission of incoming packets, the arbiter works in close co-operation with the port controllers on the line cards. Each port controller sends a regular update of its VOQ state to the arbiter. The arbiter keeps a copy of the actual state-

information of the VOQs of all connected port controllers. Based on the VOQ state-information received from the connected port controllers, the arbiter calculates the I/O configuration of the matrices and sends the result to the matrices, respectively, and the port controllers at regular intervals.

5 In Fig. 1, a known system is depicted. A plurality of line cards 1-N is connected to a number of switch cards 10. The line cards 1-N communicate with the switch cards 10 by using port controllers 1a-Na. The port controllers 1a-Na send data cells to and receive data cells from the switch cards 10 via connection lines 2, 4. The state information of the output queues of the port controllers 1a-Na are communicated to arbiters 10b via
10 communication lines 6, 8. The arbiters 10b decide which line cards 1-N are connected with each other via the lines 4, 2 to transmit respective cells in the output queues of the port controllers 1a-Na.

The transmission of cells between port controllers 1a-Na is switched by setting a crossbar matrix 10a appropriately. Input ports of the switch cards 10 are represented by
15 lines in the crossbar matrix 10a. Output port of the switch cards 10a represented by columns in the crossbar matrix 10a. To connect, for example, input port "1" with output port "3", a switch located at line 1 at column 3 of a matrix 10a is set "on".

In case the crossbar matrix 10a, or the switch 10 does not comprise a buffer memory, instant switching has to be provided, otherwise incoming packets would be
20 corrupted when switching between a first configuration to a second configuration of crossbar matrix 10a is carried out during transmission of said cell.

A crucial function for a buffered matrix to work properly is the alignment of incoming cells. Only if incoming cells are aligned, configurations of the matrix can be changed without disturbing cell transfers. An exact alignment is best for synchronisation.
25 Synchronisation of the cells at the input of the crossbar matrix is complicated, because the cables connecting port controller 1a-Na with crossbar matrix 10a may be of different lengths. Different lengths cause different signal time delay on the lines. With data rates of several Gbit/s, a difference of some centimetres of the line length results in a misalignment of the cells for a couple of bit clocks.

30 From JP 7-79218 it is known to provide a synchronisation pattern detection circuit. The synchronisation pattern detection circuit detects frames and synchronisation patterns from information stored in a shift register. By transmitting a sync signal, a fixed difference of transmission times is calculated and a synchronisation pattern is evaluated. By applying the synchronisation pattern, wire length can be different, as the input is

synchronised by using the synchronisation pattern. A problem of generating synchronising patterns is that each line card must generate such a synchronisation pattern and each switch card must derive the synchronisation pattern from the line card.

5

Thus, it is an object of the invention to provide a method and a system that allows independent alignment of cells by the line cards. A further object of the invention is to provide a flexible system configuration, being able to allow synchronisation in different configurations. A further object of the invention is to provide packet switching without buffer
10 memory in the crossbar matrices.

The objects of the invention are solved by oscillating said configuration during a set-up period between a loopback configuration and a no-transmission configuration, whereby received cells are transferred back to said sending input/output means in said loopback configuration and received cells are not transferred back to said sending
15 input/output means in said no-transmission configuration, receiving back transferred cells in said input/output means, checking said received cells in said input/output means for a transmission error, and shifting an offset of said start of cell times in case a transmission error occurred, until transferring back at least one cell is wholly carried out within a cell transfer period.

20

The configuration of the cross-connection means is changed between loopback configuration and no-transmission configuration. In loopback configuration, incoming cells are transferred back to the sending input/output means, for instance a port controller. In no-transmission configuration, incoming cells are not sent back to the sending input/output means.

25

In case cells are transferred back, these cells are received in said input/output means. As the configuration of the cross-connection means is switched between loopback configuration and no-transmission configuration, transmitted cells may be corrupted. Corruption of cells may appear, if they are received in the cross-connection means during times of configuration, e.g. cross-connection configuration periods, also called switch-over
30 periods. Corruption may also occur, if parts of the cells are received during transmission periods and parts of the cells are received during no-transmission periods. Only in case the cell is received and re-transmitted wholly within a transmission, or loopback period, it is received without transmission error.

By shifting the offset of start of cell times, the input/output means tries to find the correct time, at which cells may be transmitted and further processed within the cross-connection means without falling into a switch-over period or a no-transmission period. The offset is shifted until at least one cell is wholly re-transferred within a cell transfer period. In that case an input/output means has determined the times at which it may send cells, so that they are not corrupted due to matrix configurations.

Port controllers and separate line cards, being connected to the crossbar matrix by cables of different lengths may be aligned so that their cells are received within the matrix at equal times. The loopback configuration is active for one cell transfer period and the no-transmission configuration is also active during one cell transmission period. In no-transmission configuration, ongoing cell transfers are disturbed. Only if the offset is such that the whole cell is transmitted within loopback configuration, the input/output means are synchronised with the cross-connection means.

An alignment of all incoming cells from connected input/output means is made available according to claim 2. In this case the transmission data rate may be at its maximum, as transmission periods only last for exactly one transmission of a cell.

A central clock signal, according to claim 3, allows easy cell synchronisation.

Serialisation and de-serialisation, according to claim 4 and 5, allows serialised transmission of data packets. Evaluating a bit error indicator, according to claim 6, may also be carried out based on a coding scheme applied for transmission on a transmission line. The output of the bit error indication may be used as decision to change the offset of start of cell times in order to delay the output of cells.

Using an offset counter, according to claim 7, allows a shift of the start of cell times of outgoing cells with respect to a matrix configuration synchronisation signal. The delay is incremented or decremented by defined step widths.

The offset counter may also be controlled, according to claim 8, in a way that the start of cell time is pre-running with respect to the start of cell signal as much as possible without generation of bit errors, and afterwards delaying the start of cell time as much as possible until bit errors occur. The length of cell transmission periods may then be adjusted and the total throughput is increased. The bit error rate is brought to a minimum in case the start of cell time is set according to claim 9.

According to a further aspect of the invention, a packet switch is provided, where said port controller comprises a start of cell signal generator for generating start of cell signals, an offset controller for shifting a start of cell time based on said start of cell signal,

and an error detection mean for detecting corrupt received cells and where said cross-connection means comprises a configuration controller for controlling an oscillation between a loopback configuration and a no-transmission configuration of said cross-connection means.

5 A packet switch, according to claim 11, is advantageous, as a central clock signal allows exact synchronisation between port controllers and cross-connection means.

By providing serialisation and de-serialisation means, according to claim 12, data bits or packets may be serially transferred.

10 Providing an NxN cross matrix, according to claims 13 and 14, configuration changes may be applied easily. In such a matrix, a line corresponds to an input port and a column corresponds to an output port. A switch at position (X, Y) in said NxN matrix connects input port X with output port Y.

By providing a bit error indicator, according to claim 15, a bit error may be derived from the coding scheme applied for a transmission on a line.

15 Another aspect of the invention is the use of a described method or a described packet switch in packet switched networks for synchronising start of cell times in various port controllers during a set-up, to allow configuration changes in cross-connection means without disturbing cell transfers.

20 These and other aspects of the invention will be apparent from and elucidated with reference to the embodiments described hereinafter.

In the figures show:

Fig. 1 a packet switch configuration;

25 Fig. 2 diagrammatically the variation of the offset of start of cell times;

Fig. 3 a block diagram of an inventive packet switch.

30 Fig. 2 depicts in diagrams A-E a succession of cells being send within a set-up period of a packet switch.

As can be seen in diagram A, data cells 12 are usually sent at start of cell times 14, which usually correspond to a start of cell signal 16. Usually, a data cell 12 is sent exactly after a start of cell signal 16 is generated. In that case, all cells from all port controllers are sent at the same time. As cables connecting the port controllers with the crossbar matrices

may have different length, the cells are received in the matrices at different times. For switching packet switched connections, the configuration of the matrices have to be changed. During these configuration periods 20, cells 12 may be received within the matrices, which causes cell corruption. The invention provides a method to set up a packet switch in a way, that cell corruption may be avoided.

The configuration of a crossbar matrix of a switch card is changed, as depicted in diagram B, between configuration B_1 and configuration B_0 during a set-up period.

Configuration B_1 stands for loopback configuration, where the matrix is a unit matrix, and configuration B_0 stands for no-transmission configuration, where the matrix is a null matrix.

During a set-up period of the crossbar matrix, the configuration is changed between B_1 , and B_0 . In configuration B_1 , cells 12 sent from a port controller 1a are sent back to this respective port controller 1a. During no-transmission time B_0 , cells 12 sent from port controller 1a are not sent back to port controller 1a. The inventive method works as follows.

As depicted in diagram C, start of cell signals 16 are generated in all port controllers according to a common clock signal at system start. Cells 12 are sent at start of cell times 14 and received in the crossbar matrix during loopback configuration B_1 , no-transmission configuration B_0 , and set-up period 20.

First, cells are released from the cell memory at start of cell times 14 directly after reception of the start of cell signal 16. An offset may be generated, by which the start of cell time 14 is shifted from the start of cell signal 16 by a certain value. In case sent cells 12 are received in the matrix during set-up period 20 or no-transmission configuration B_0 , they are not correctly sent back to the sending port controller. The received cells are corrupt. To recognise corrupt cells, each received cell is evaluated and thus transmission errors are detected. In case a transmission error is detected, the cell 12 has been received in the matrix at times where the transmission becomes corrupt, e.g. no-transmission period B_0 or set-up period 20.

The offset is increased incrementally, the start of cell time 14 is incrementally delayed from the start of cell signal 16, until the cell 12 will be sent at a time where it can be received in the matrix within the loopback configuration time B_0 and thus sent back to the port controller without error. The offset will then not be increased further, and may be used during the operation of port controller, as it allows an error-free transmission of cells.

In diagram D, the offset of the start of cell time 14 is decreased. First, a cell is sent at times at which it is received in the matrix during set up period 20. In that case a cell 12 is corrupted. By shifting the offset, the start of cell time 14 is pre-running the start of cell

signal to a further extend. After a few shifts, the start of cell time 14 will be such, that a cell 12 will be wholly received within said crossbar matrix during period B₁.

According to diagram E, the offset is controlled in a way that the start of cell time 14 is at first pre-running with respect to the start of cell signal 16 as much as possible without generating bit errors, and afterwards decreased as much as possible until bit errors start appearing. By this, the duration of a loopback configuration B₁ is evaluated and the start of cell time 14 can be adjusted such that a cell 12 will be transmitted in the middle of transmission time. Also the gap 13 between cells 12 may thus be decreased to a minimum value.

Fig. 3 depicts a port controller 1 and a crossbar matrix 40. Port controller 1 provides an input port 24, an output port 26, a cell memory 28, a start of cell signal generator 30, an offset counter 32, a serialiser 34, a de-serialiser 36 and a bit error indicator 38. Crossbar matrix 40 provides cell input ports 41, cell output port 43, switched connections 42 and interrupted connected 44. Further depicted is a central clock generator 48 and a configuration controller 46.

Incoming packets at port 24 are segmented into fixed sizes packet fragments, cells, and stored in cell memory 28. Outgoing cells are reassembled back into packets and put out at port 26. During set-up of a port controller 1, cells are sent at start of cell times which are generated at the start of cell time generator. These packets are serialised in serialiser 34 and sent to cell input port 41a. In loopback configuration, cell input port 41a is switched to cell output port 43a.

Loopback configuration means that a unit matrix of the switches 42, 44 is generated, whereby an input port 41a, b is switched to an output port 43a, b, respectively. The unit matrix means that each input port is switched to its respective output port and no other connection is switched. These ports being switched belong to one same port controller. The switches 42 are "on" the switches 44 are "off". This is the so-called loopback configuration.

No-transmission configuration is realised by a null matrix, where no connections between input ports and output ports are switched at all. All incoming data-packets are lost or corrupted. All switches 42, 44 are "off".

During set-up the crossbar matrix 40 is switched between unit matrix and null matrix by configuration controller 46, which is controlled by a system clock signal generated by central clock generator 48. The system clock generated by central clock generator 48 is also provided to offset counter 32 and to start of cell signal generator 30. Cells which are

retransmitted to port controller 1 are received in the serialiser 36. The received cells are evaluated in bit error rate indicator 38. In case a bit error occurred, the offset counter 32 is increased. By increasing the offset counter 32 the start of cell signal generator generates a start of cell time prevailing the central clock signal by the amount of the offset counter 32. By increasing the offset counter 32, the start of cell time will be changed until a cell is transmitted to crossbar matrix 40 and received by de-serialiser 36 without transmission errors, which means that the cell is received in crossbar matrix 40 in a transmission period.

By applying the offset counter and the bit error indicator 38, start of cell times may be synchronised so as to align incoming cells at matrix 40 from various port controllers

1.

No central synchronisation mechanism or line length measurement is needed, rather all port controllers adjust the cell alignment autonomously.

Reference Signs

	1, N	line card
	1a, Na	port controller
5	2, 4	transmission connection
	6, 8	signalling connection
	10	switch card
	10a	crossbar matrix
	10b	arbiter
10	12	cell
	13	transmission gap
	14	start of cell time
	16	start of cell signal
	B ₁	loopback configuration
15	B ₀	no-transmission configuration
	20	set-up period
	24	input
	26	output
	28	cell memory
20	30	start of cell signal generator
	32	offset counter
	34	serialiser
	36	de-serialiser
	38	bit error indicator
25	40	crossbar matrix
	41	cell input port
	42	switched connection
	43	cell output port
	44	interrupted connection
30	46	configuration controller
	48	central clock generator